#### <sup>b</sup> UNIVERSITÄT BERN

 $u^{\scriptscriptstyle b}$ 

### **Virtual Machines**

#### **Guest lecture – Adrian Lienhard**



#### **Birds-eye view**



A virtual machine is an abstract computing architecture supporting a programming language in a hardware-independent fashion



Z1, 1938

### Roadmap



- > Introduction
- > The heap store
- > Interpreter
- > Automatic memory management
- > Threading System
- > Optimizations

### **Implementing a Programming Language**

#### Compilation



#### Interpretation



#### **Compilation + Interpretation**



#### How are VMs implemented?

Typically using an *efficient and portable language* such as C, C++, or assembly code

Pharo VM platform-independent part written in *Slang:* 

- subset of Smalltalk, translated to C
- core: 600 methods or 8k LOC in Slang
- Slang allows one to simulate VM in Smalltalk

#### Main Components of a VM



The heap store Interpreter Automatic memory management Threading System

### **Pros and Cons of the VM Approach**

#### Pros

- > Platform independence of application code "Write once, run anywhere"
- > Simpler programming model
- > Security
- > Optimizations for different hardware architectures

#### Cons

- > Execution overhead
- > Not suitable for system programming

### Roadmap



- > Introduction
- > The heap store
- > Interpreter
- > Automatic memory management
- > Threading System
- > Optimizations

# **Object Memory Layout**



# **Different Object Formats**

- > fixed pointer fields
- > indexable types:
  - indexable pointer fields (e.g., Array)
  - indexable weak pointer fields (e.g., WeakArray)
  - indexable word fields (e.g., Bitmap)
  - indexable byte fields (e.g., ByteString)

Object format (4bit)

- no fields 0
- fixed fields only
- indexable pointer fields only
- 2 3 both fixed and indexable pointer fields
- 4 both fixed and indexable weak fields
- indexable word fields only 6
- 8-11 indexable byte fields only
- 12-15 ....

#### **Iterating Over All Objects in Memory**

"Answer the first object on the heap" anObject someObject

"Answer the next object on the heap" anObject nextObject

Excludes small integers!

### Roadmap



- > Introduction
- > The heap store
- > Interpreter
- > Automatic memory management
- > Threading System
- > Optimizations

#### **Stack vs. Register VMs**

VM provides a virtual processor that interprets bytecode instructions

Stack machines

- Smalltalk, Java and most other VMs
- Simple to implement for different hardware architectures

**Register machines** 

- only few register VMs, e.g., Parrot VM (Perl6)
- potentially faster than stack machines

#### **Interpreter State and Loop**

#### Interpreter state

- instruction pointer (ip): points to current bytecode
- stack pointer (sp): topmost item in the operand stack
- current active method or block context
- current active receiver and method

Interpreter loop

- 1. branch to appropriate bytecode routine
- 2. fetch next bytecode
- 3. increment instruction pointer
- 4. execute the bytecode routine
- 5. return to 1.

#### **Method Contexts**

method context



method header:

- primitive index
- number of args
- number of temps
- large context flag
- number of literals

#### **Stack Manipulating Bytecode Routine**

#### Example: bytecode <70> self

Interpreter>>pushReceiverBytecode
 self fetchNextBytecode.
 self push: receiver

Interpreter>>push: anObject
 sp := sp + BytesPerWord.
 self longAt: sp put: anObject

### **Stack Manipulating Bytecode Routine**

#### Example: bytecode <01> pushRcvr: 1

```
Interpreter>>pushReceiverVariableBytecode
  self fetchNextBytecode.
  self pushReceiverVariable: (currentBytecode bitAnd: 16rF)
Interpreter>>pushReceiverVariable: fieldIndex
  self push: (
    self fetchPointer: fieldIndex ofObject: receiver)
Interpreter>>fetchPointer: fieldIndex ofObject: oop
  ^ self longAt: oop + BaseHeaderSize + (fieldIndex * BytesPerWord)
```

### **Message Sending Bytecode Routine**

Example: bytecode <E0> send: hello

- 1. find selector, receiver and its class
- 2. lookup message in the class' method dictionary
- 3. if method not found, repeat this lookup in successive superclasses; if superclass is nil, instead send #doesNotUnderstand:
- 4. create a new method context and set it up
- 5. activate the context and start executing the instructions in the new method

#### **Message Sending Bytecode Routine**

#### Example: bytecode <E0> send: hello

```
Interpreter>>sendLiteralSelectorBytecode
  selector := self literal: (currentBytcode bitAnd: 16rF).
  argumentCount := ((currentBytecode >> 4) bitAnd: 3) - 1.
  rcvr := self stackValue: argumentCount.
  class := self fetchClassOf: rcvr.
  self findNewMethod.
  self executeNewMethod.
  self fetchNewBytecode
```

This routine (bytecodes 208-255) can use any of the first 16 literals and pass up to 2 arguments

```
E0(hex) = 224(dec)
= 1110 0000(bin)
E0 AND F = 0
=> literal frame at 0
((E0 >> 4) AND 3) - 1 = 1
=> 1 argument
```

#### **Primitives**

Primitive methods trigger a VM routine and are executed without a new method context unless they fail

ProtoObject>>nextObject
<primitive: 139>
self primitiveFailed

- > Improve performance (arithmetics, at:, at:put:, ...)
- Do work that can only be done in VM (new object creation, process manipulation, become, ...)
- > Interface with outside world (keyboard input, networking, ...)
- > Interact with VM plugins (named primitives)

### Roadmap



- > Introduction
- > The heap store
- > Interpreter
- > Automatic memory management
- > Threading System
- > Optimizations

#### **Automatic Memory Management**

Tell when an object is no longer used and then recycle the memory



- Challenges
- Fast allocation
- Fast program execution
- Small predictable pauses
- Scalable to large heaps
- Minimal space usage

# **Main Approaches**

1. Reference Counting

2. Mark and Sweep

# **Reference Counting GC**

#### Idea

- > For each store operation increment count field in header of newly stored object
- > Decrement if object is overwritten
- If count is 0, collect object and decrement the counter of each object it pointed to

#### Problems

- > Run-time overhead of counting (particularly on stack)
- > Inability to detect cycles (need additional GC technique)

# **Reference Counting GC**



# Mark and Sweep GC

#### Idea

- > Suspend current process
- Mark phase: trace each accessible object leaving a mark in the object header (start at known root objects)
- > Sweep phase: all objects with no mark are collected
- > Remove all marks and resume current process

#### Problems

- > Need to "stop the world"
- > Slow for large heaps →generational collectors
- > Fragmentation → compacting collectors

### Mark and Sweep GC



# **Generational Collectors**

Most new objects live very short lives, most older objects live forever [Ungar 87]

#### Idea

- > Partition objects in generations
- > Create objects in young generation
- > Tenuring: move live objects from young to old generation
- > Incremental GC: frequently collect young generation (very fast)
- > Full GC: infrequently collect young+old generation (slow)

#### Difficulty

> Need to track pointers from old to new space

### **Generational Collectors: Remembered Set**

Write barrier: remember objects with old-young pointers:

 On each store check whether storee (object2) is young and storand (object1) is old

object1.f := object2

- > If true, add storand to remembered set
- > When marking young generation, use objects in remembered set as additional roots



# **Compacting Collectors**

Idea

- > During the sweep phase all live objects are packed to the beginning of the heap
- > Simplifies allocation since free space is in one contiguous block

#### Challenge

- > Adjust all pointers of moved objects
  - object references on the heap
  - pointer variables of the interpreter!

#### **The Pharo GC**

Pharo: mark and sweep compacting collector with two generations

- Cooperative, i.e., not concurrent
- Single threaded

#### When Does the GC Run?

- Incremental GC on allocation count or memory needs
- Full GC on memory needs
- Tenure objects if survivor threshold exceeded



#### **VM Memory Statistics**

SmalltalkImage current
 vmStatisticsReportString

memory	20,245,0	028 bytes
	old	14,784,388 bytes (73.0%)
	young	117,724 bytes (0.6%)
	used	14,902,112 bytes (73.6%)
	free	5,342,916 bytes (26.4%)
GCs		975 (48ms between GCs)
	full	0 totalling 0ms (0.0% uptime)
	incr	975 totalling 267ms (1.0% uptime), avg 0.0ms
	tenures 14 (avg	69 GCs/tenure)
Since 1	last view	90 (54ms between GCs)
	uptime	4.8s
	full	0 totalling 0ms (0.0% uptime)
	incr	90 totalling 29ms (1.0% uptime), avg 0.0ms
	tenures 1 (avg 90 GCs/tenure)	

#### **Memory System API**

"Force GC" Smalltalk garbageCollectMost Smalltalk garbageCollect

"Is object in remembered set, is it young?" Smalltalk rootTable includes: anObject Smalltalk isYoung: anObject

"Various settings and statistics" SmalltalkImage current getVMParameters

"Do an incremental GC after this many allocations" SmalltalkImage current vmParameterAt: 5 put: 4000. "Tenure when more than this many objects survive the GC" SmalltalkImage current vmParameterAt: 6 put: 2000. "Grow/shrink headroom" SmalltalkImage current vmParameterAt: 25 put: 4\*1024\*1024. SmalltalkImage current vmParameterAt: 24 put: 8\*1024\*1024.

### **Finding Memory Leaks**

I have objects that do not get collected. What's wrong?

- maybe object is just not GCed *yet* (force a full GC!)
- find the objects and then explore who references them

# PointerFinder finds a path from a root to some object

PointerFinder on: AssignmentNode someInstance

PointerExplorer new openExplorerFor: AssignmentNode someInstance



### Roadmap



- > Introduction
- > The heap store
- > Interpreter
- > Automatic memory management
- > Threading System
- > Optimizations

### **Threading System**

Multithreading is the ability to create concurrently running "processes"

#### Non-native threads (green threads)

- Only one native thread used by the VM
- Simpler to implement and easier to port

#### Native threads

- Using the native thread system provided by the OS
- Potentially higher performance

#### **Pharo: Green Threads**

Each process has its own execution stack, ip, sp, ... There is always one (and only one) running process Each process behaves as if it owns the entire VM Each process can be interrupted (→context switching)

#### **Representing Processes and Run Queues**





### **Context Switching**

Interpreter>>transferTo: newProcess

- 1. store the current ip and sp registers to the current context
- 2. store the current context in the old process' suspendedContext
- 3. change Processor to point to newProcess
- 4. load ip and sp registers from new process' suspendedContext

When you perform a context switch, which process should run next?

#### **Process Scheduler**

- > Cooperative between processes of the same priority
- > *Preemptive* between processes of different priorities

Context is switched to the first process with highest priority when:

- current process waits on a semaphore
- current process is suspended or terminated
- Processor yield is sent

Context is switched if the following process has a higher priority:

- process is resumed or created by another process
- process is resumed from a signaled semaphore

When a process is interrupted, it moves to the back of its run queue

#### **Example: Semaphores and Scheduling**

```
here := false.
lock := Semaphore forMutualExclusion.
[lock critical: [here := true]] fork.
lock critical: [
  self assert: here not.
Processor yield.
self assert: here not].
Processor yield.
self assert: here
```

### Roadmap



- > Introduction
- > The heap store
- > Interpreter
- > Automatic memory management
- > Threading System
- > **Optimizations**

### Many Optimizations...

- > Method cache for faster lookup: receiver's class + method selector
- Method context cache (as much as 80% of objects created are context objects!)
- > Interpreter loop: 256 way case statement to dispatch bytecodes
- > Quick returns: methods that simply return a variable or known constant are compiled as a primitive method
- Small integers are tagged pointers: value is directly encoded in field references. Pointer is tagged with low-order bit equal to 1. The remaining 31 bit encode the signed integer value.
- > ...

# **Optimization: JIT (not in Pharo)**

#### Idea of Just In Time Compilation

- > Translate unit (method, loop, ...) into native machine code at runtime
- > Store native code in a buffer on the heap

#### Challenges

- > Run-time overhead of compilation
- > Machine code takes a lot of space (4-8x compared to bytecode)
- > Deoptimization is very tricky

Adaptive compilation: gather statistics to compile only units that are heavily used (*hot spots*)

#### References

> Virtual Machines, Iain D. Craig, Springer, 2006

> Back to the Future – The Story of Squeak, A Practical Smalltalk Written in Itself, Ingalls, Kaehler, Maloney, Wallace, Kay, OOPSLA '97

> Smalltalk-80, the Language and Its Implementation (the Blue Book), Goldberg, Robson, Addison-Wesley, '83 http://stephane.ducasse.free.fr/ FreeBooks/BlueBook/Bluebook.pdf

> The Java Virtual Machine Specification, Second Edition, http:// java.sun.com/docs/books/jvms/

> Stacking them up: a Comparison of Virtual Machines, Gough, IEEE'01

> Virtual Machine Showdown: Stack Versus Registers, Shi, Gregg, Beatty, Ertl, VEE'05

# What you should know!

- What is the difference between the operand stack and the execution stack?
- How do bytecode routines and primitives differ?
- Why is the object format encoded in a complicated 4bit pattern instead of using regular boolean values?
- Why is the object address not suitable as a hash value?
- What happens if an object is only weakly referenced?
- Why is it hard to build a concurrent mark sweep GC?
- What does *cooperative multithreading* mean?
- How do you protect code from concurrent execution?

#### Can you answer these questions?

- There is a lot of similarity between VM and OS design. What are the common components?
- Why is accessing the 16th instance variable of an object more efficient than the 17th?
- Which disastrous situation could occur if a local C pointer variable exists when a new object is allocated?
- Why does #allObjectsDo: not include small integers?
- What is the largest possible small integer?

#### License

#### http://creativecommons.org/licenses/by-sa/3.0/



#### **Attribution-ShareAlike 3.0 Unported**

You are free:

to Share — to copy, distribute and transmit the work

to Remix — to adapt the work

#### Under the following conditions:

**Attribution.** You must attribute the work in the manner specified by the author or licensor (but not in any way that suggests that they endorse you or your use of the work).

**Share Alike.** If you alter, transform, or build upon this work, you may distribute the resulting work only under the same, similar or a compatible license.

For any reuse or distribution, you must make clear to others the license terms of this work. The best way to do this is with a link to this web page.

Any of the above conditions can be waived if you get permission from the copyright holder. Nothing in this license impairs or restricts the author's moral rights.